

[0002] Inventors:  
Isaac Levanon  
Yonathan Lavi

[0004] Cross-Reference to Related Applications

[0006] Background of the Invention

[0008] The present invention is related to network based, image distribution systems and, in particular, to a system and methods for efficiently

selecting and distributing image parcels through a narrowband or otherwise limited bandwidth communications channel to support presentation of high-resolution images subject to dynamic viewing frustums.

[0009]        Description of the Related Art:

[0010]        The Internet and other network systems provide a unique opportunity to transmit complex images, typically large scale bit-maps, particularly those approaching photo-realistic levels, over large distances. In common application, the images are geographic, topographic, and other highly detailed maps. The data storage requirements and often proprietary nature of such images are such that conventional interests are to transfer the images on an as-needed basis.

[0011]        In conventional fixed-site applications, the image data is transferred over a relatively high-bandwidth network to client computer systems that, in turn, render the image. Client systems typically implement a local image navigation system to provide zoom and pan functions based on user interaction. As well recognized problem with such conventional systems is that full resolution image presentation is subject to the inherent transfer latency of the network. Different conventional systems have been proposed to reduce the latency affect by transmitting the image in highly compressed formats that support progressive resolution build-up of the image within the current client field of view. Using a transform compressed image transfer function increases the field of the image that can be transferred over a fixed bandwidth network in unit time. Progressive image resolution transmission, typically using a differential resolution method,

permits an approximate image to be quickly presented with image details being continuously added over time.

[0012] Tzou, in U.S. Patent Number 4,698,689, describes a two-dimensional data transform system that supports transmission of differential coefficients to represent an image. Subsequent transmitted coefficient sets are progressively accumulated with prior transmitted sets to provide a succeeding refined image. The inverse-transform function performed by the client computer is, however, highly compute intensive. In order to simplify the transform implementation and further reduce the latency of presenting any portion of an approximate image, images are sub-divided into a regular array. This enables the inverse-transform function on the client, which is time-critical, to deal with substantially smaller coefficient data sets. The array size in Tzou is fixed, which leads to progressively larger coefficient data sets as the detail level of the image increases. Consequently, there is an inherently increasing latency in resolving finer levels of detail.

[0013] An image visualization system proposed by Yap et al., U.S. Patent Number 6,182,114, overcomes some of the foregoing problems. The Yap et al. system also employs a progressive encoding transform to compress the image transfer stream. The transform also operates on a subdivided image, but the division is indexed to the encoding level of the transform. The encoded transform coefficient data sets are, therefore, of constant size, which supports a modest improvement in the algorithmic performance of the inverse transform operation required on the client.

[0014] Yap et al. adds utilization of client image panning or other image pointing input information to support a foveation-based operator to influence the

retrieval order of the subdivided image blocks. This two-dimensional navigation information is used to identify a foveal region that is presumed to be the gaze point of a client system user. The foveation operator defines the corresponding image block as the center point of an ordered retrieval of coefficient sets representing a variable resolution image. The gaze point image block represents the area of highest image resolution, with resolution reduction as a function of distance from the gaze point determined by the foveation operator. This technique thus progressively builds image resolution at the gaze point and succeeding outward based on a relatively compute intensive function. Shifts in the gaze point can be responded to with relative speed by preferentially retrieving coefficient sets at and near the new foveal region.

[0015] Significant problems remain in permitting the convenient and effective use of complex images by many different types of client systems, even with the improvements provided by the various conventional systems. In particular, the implementation of conventional image visualization systems is generally unworkable for smaller, often dedicated or embedded, clients where use of image visualization would clearly be beneficial. Conventional approaches effectively presume that client systems have an excess of computing performance, memory and storage. Small clients, however, typically have restricted performance processors with no dedicated floating-point support, little general purpose memory, and extremely limited persistent storage capabilities, particularly relative to common image sizes. A personal digital assistant (PDA) is a characteristic small client. Embedded, low-cost kiosk and automobile navigation systems are other typical examples. Such systems are not readily capable, if at all,

of performing complex, compute-intensive Fourier or wavelet transforms, particularly within a highly restricted memory address space.

[0016] As a consequence of the presumption that the client is a substantial computing system, conventional image visualization systems also presume that the client is supported by a complete operating system. Indeed, many expect and require an extensive set of graphics abstraction layers to be provided by the client system to support the presentation of the delivered image data. In general, these abstraction layers are conventionally considered required to handle the mapping of the image data resolution to the display resolution capabilities of the client system. That is, resolution resolved image data provided to the client is unconstrained by any limitation in the client system to actually display the corresponding image. Consequently, substantial processor performance and memory can be conventionally devoted to handling image data that is not or cannot be displayed.

[0017] Another problem is that small clients are generally constrained to generally to very limited network bandwidths, particularly when operating under wireless conditions. Such limited bandwidth conditions may exist due to either the direct technological constraints dictated by the use of a low bandwidth data channel or indirect constraints imposed on relatively high-bandwidth channels by high concurrent user loads. Cellular connected PDAs and webphones are examples of small clients that are frequently constrained by limited bandwidth conditions. The conventionally realizable maximum network transmission bandwidth for such small devices may range from below one kilobit per second to several tens of kilobits per second. While Yap et al. states that the described system can work over low bandwidth lines, little more than utilizing wavelet-based

data compression is advanced as permitting effective operation at low communications bandwidths. While reducing the amount of data that must be carried from the server to the client is significant, Yap et al. simply relies on the data packet transfer protocols to provide for an efficient transfer of the compressed image data. Reliable transport protocols, however, merely mask packet losses and the resultant, sometimes extended, recovery latencies. When such covered errors occur, however, the aggregate bandwidth of the connection is reduced and the client system can stall waiting for further image data to process.

[0018] Consequently, there remains a need for an image visualization system that can support small client systems, place few requirements on the supporting client hardware and software resources, and efficiently utilize low to very low bandwidth network connections.

[0019] Summary of the Invention

[0020] Thus, a general purpose of the present invention is to provide an efficient system and methods of optimally presenting image data on client systems with potentially limited processing performance, resources, and communications bandwidth.

[0021] This is achieved in the present invention by providing for the dynamic visualization of image data provided through a network communications channel by a client system including a parcel request subsystem and a parcel rendering subsystem. The parcel request subsystem includes a parcel request queue and is operative to request discrete image data parcels in a priority order and to store received image data parcels in a parcel data store. The parcel

request subsystem is responsive to an image parcel request of assigned priority to place the image parcel request in the parcel request queue ordered in correspondence with the assigned priority. The parcel rendering subsystem is coupled to the parcel data store to selectively retrieve and render received image data parcels to a display memory. The parcel rendering system provides the parcel request subsystem with the image parcel request of the assigned priority.

[0022] An advantage of the present invention is that both image parcel data requests and the rendering of image data are optimized to address the display based on the display resolution of the client system.

[0023] Another advantage of the present invention is that the prioritization of image parcel requests is based on an adaptable parameter that minimizes the computational complexity of determining request prioritization and, in turn, the progressive improvement in display resolution within the field of view presented on a client display.

[0024] A further advantage of the present invention is that the client software system requires relatively minimal client processing power and storage capacity. Compute intensive numerical calculations are minimally required and image parcel data is compactly stored in efficient data structures. The client software system is very small and easily downloaded to conventional computer systems or embedded in conventional dedicated function devices, including portable devices, such as PDAs and webphones.

[0025] Still another advantage of the present invention is that image parcel data requests and presentation can be readily optimized to use low to very low bandwidth network connections. The software system of the present invention provides for re-prioritization of image parcel data requests and presentation in

circumstances where the rate of point-of-view navigation exceeds the data request rate.

[0026] Yet another advantage of the present invention is that image parcel data rendering is performed without requiring any complex underlying hardware or software display subsystem. The client software system of the present invention includes a bit-map rendering engine that draws directly to the video memory of the display, thus placing minimal requirements on any underlying embedded or disk operating system and display drivers. Complex graphics and animation abstraction layers are not required.

[0027] Still another advantage of the present invention is that image parcel block compression is used to obtain fixed size transmission data blocks. Image parcel data is recoverable from transmission data using a relatively simple client decompression algorithm. Using fixed size transmission data blocks enables image data parcels to be delivered to the client in bounded time frames.

[0028] A yet further advantage of the present invention is that multiple data forms can be transferred to the client software system for concurrent display. Sparse array overlay data, correlated positionally to the image parcel data and generally insensitive to image parcel resolution, can be initially or progressively provided to the client for parsing and parallel presentation on a client display image view.

[0029] Brief Description of the Drawings

[0030] These and other advantages and features of the present invention will become better understood upon consideration of the following detailed description of the invention when considered in connection with the accompanying



drawings, in which like reference numerals designate like parts throughout the figures thereof, and wherein:

[0031] Figure 1 depicts a preferred system environment within which various embodiments of the present invention can be utilized;

[0032] Figure 2 is a block diagram illustrating the preparation of image parcel and overlay data set that are to be stored by and served from a network server system in accordance with a preferred embodiment of the present invention;

[0033] Figure 3 is a block diagram of a client system image presentation system constructed in accordance with a preferred embodiment of the present invention;

[0034] Figure 4 provides a data block diagram illustrating an optimized client image block processing path constructed in accordance with a preferred embodiment of the present invention;

[0035] Figure 5 is a process flow diagram showing a main processing thread implemented in a preferred embodiment of the present invention;

[0036] Figure 6 provides a process flow diagram showing a network request thread implemented in a preferred embodiment of the present invention;

[0037] Figure 7 provides a process flow diagram showing a display image rendering thread implemented in a preferred embodiment of the present invention;

[0038] Figure 8 provides a process flow diagram showing the parcel map processing performed preliminary to the rendering of image data parcels in accordance with a preferred embodiment of the present invention;

[0039] Figure 9 provides a process flow diagram detailing the rendering and progressive prioritization of image parcel data download requests in accordance with a preferred embodiment of the present invention; and

[0040] Figure 10 provides a process flow diagram detailing the determination of an optimal detail level for image parcel presentation for a current viewing frustum in accordance with a preferred embodiment of the present invention.

[0041] Detailed Description of the Invention

[0042] The preferred operational environment 10 of the present invention is generally shown in Figure 1. A network server system 12, operating as a data store and server of image data, is responsive to requests received through a communications network, such as the Internet 14 generally and various tiers of internet service providers (ISPs) including a wireless connectivity provider 16. Client systems, including conventional workstations and personal computers 18 and smaller, typically dedicated function devices often linked through wireless network connections, such as PDAs, webphones 20, and automobile navigation systems, source image requests to the network server 12, provide a client display and enable image navigational input by a user of the client system. Alternately, a dedicated function client system 20 may be connected through a separate or plug-in local network server 22, preferably implementing a small, embedded Web server, to a fixed or removable storage local image repository 24. Characteristically, the client system 18, 20 displays are operated at some fixed resolution generally dependent on the underlying display hardware of the client systems 18, 20.

[0043] The image navigation capability supported by the present invention encompasses a viewing frustum placed within a three-dimensional space over the imaged displayed on the client 18, 20. Client user navigational inputs are supported to control the x, y lateral, rotational and z height positioning of the viewing frustum over the image as well as the camera angle of incidence relative to the plane of the image. To effect these controls, the software implemented on the client systems 18, 20 supports a three-dimensional transform of the image data provided from the server 12, 22.

[0044] In accordance with the preferred embodiments of the present invention, as generally illustrated in Figure 2, a network image server system 30 stores a combination of source image data 32 and source overlay data 34. The source image data 32 is typically high-resolution bit-map satellite imagery of geographic regions, which can be obtained from commercial suppliers. The overlay image data 34 is typically a discrete data file providing image annotation information at defined coordinates relative to the source image data 32. In the preferred embodiments of the present invention, image annotations include, for example, street, building and landmark names, as well as representative 2 and 3D objects, graphical icons, decals, line segments, and text and other characters.

[0045] The network image server system 30 preferably pre-processes the source image data 32 and source overlay data 34 to forms preferred for storage and serving by the network server 12, 22. The source image data 32 is preferably pre-processed to obtain a series  $K_{1-N}$  of derivative images of progressively lower image resolution. The source image data 32, corresponding to the series image  $K_0$ , is also subdivided into a regular array such that each resulting image parcel of the array has a 64 by 64 pixel resolution where the image data has a color or

bit per pixel depth of 16 bits, which represents a data parcel size of 8K bytes. The resolution of the series  $K_{1..N}$  of derivative images is preferably related to that of the source image data 32 or predecessor image in the series by a factor of four. The array subdivision is likewise related by a factor of four such that each image parcel is of a fixed 8K byte size.

[0046] In the preferred embodiment of the present invention, the image parcels are further compressed and stored by the network server 12, 22. The preferred compression algorithm implements a fixed 4:1 compression ratio such that each compressed and stored image parcel has a fixed 2K byte size. The image parcels are preferably stored in a file of defined configuration such that any image parcel can be located by specification of a  $K_D$ , X, Y value, representing the image set resolution index D and corresponding image array coordinate.

[0047] The source overlay data 34 is preferably pre-processed 36 into either an open XML format, such as the Geography Markup Language (GML), which is an XML based encoding standard for geographic information developed by the OpenGIS Consortium (OGC; [www.opengis.org](http://www.opengis.org)), or a proprietary binary representation. The XML/GML representation is preferred as permitting easier interchange between different commercial entities, while the binary representation is preferred as more compact and readily transferable to a client system 18, 20. In both cases, the source overlay data 34 is pre-processed to contain the annotation data preferably in a resolution independent form associated with a display coordinate specification relative to the source image data 32. The XML, GML or binary overlay data may be compressed prior to storage on the network server 12, 22.

[0048] The preferred architecture 40 of a client system 18, 20, for purposes of implementing the present invention, is shown in Figure 3. The architecture 40 is preferably implemented by a software plug-in or application executed by the client system 18, 20 and that utilizes basic software and hardware services provided by the client system 18, 20. A parcel request client 42 preferably implements an HTML client that supports HTML-based interactions with the server 12, 22 using the underlying network protocol stack and hardware network interface provided by the client systems 18, 20. A central parcel processing control block 44 preferably implements the client process and control algorithms. The control block 44 directs the transfer of received image parcels and XML/GML/binary overlay data to a local parcel data store 46. Preferably image data parcels are stored in conventional quad-tree data structures, where tree nodes of depth D correspond to the stored image parcels of a derivative image of resolution  $K_D$ . The XML/GML/binary overlay data is preferably stored as a data object that can be subsequently read by an XML/GML/binary parser implemented as part of the control block 44.

[0049] The control block 44 is also responsible for decompressing and directing the rendering of image parcels to a local display by a rendering engine 48. Preferably, the rendering engine 48 writes to the video memory of the underlying client display hardware relying on only generic graphics acceleration hardware capabilities. In general, the relied on capabilities include bit-blt and related bit-oriented functions that are readily supported by current conventional display controller hardware. The rendering engine 48 is optimized to perform image parcel texture mapping without reliance on complex floating point

operations, permitting even relatively simple processors to efficiently execute the rendering engine 48.

[0050] Changes in the viewing frustum are determined from user input navigation commands by a frustum navigation block 50. In the preferred embodiments of the present invention, the input navigation controls are modeled for three-dimensional fly-over navigation of the displayed image. The navigation controls support point-of-view rotation, translation, attitude, and altitude over the displayed image. The effective change in viewing frustum as determined by the frustum navigation block 50 is provided to the control block 44.

[0051] The control block 44, based in part on changes in the viewing frustum, determines the ordered priority of image parcels to be requested from the server 12, 22 to support the progressive rendering of the displayed image. The image parcel requests are placed in a request queue 52 for issuance by the parcel request client 42. Preferably, the pending requests are issued in priority order, thereby dynamically reflecting changes in the viewing frustum with minimum latency.

[0052] An optimal image parcel data flow 60, as configured for use in the preferred embodiments of the present invention, is shown in Figure 4. Preferably, the TCP/IP network protocol is used to deliver image parcels to the clients 18, 20. For the preferred embodiments, where network bandwidth is limited or very limited, entire image parcels are preferably delivered in corresponding data packets. This preference maximizes data delivery while avoiding the substantial latency and processing overhead of managing image parcel data split over multiple network packets. Thus, a 2K byte compressed image parcel 62 is delivered as the data payload of a TCP/IP packet 64. Uncompressed, the 8K

byte image parcel 62 is recognized as part of the present invention as being within the nominally smallest L1 data cache 66 size of conventional microprocessors 68. By ensuring that an uncompressed image parcel fits within the L1 cache, the texture map rendering algorithm can execute with minimum memory management overhead, thus optimally utilizing the processing capability of the microprocessor 68. Additionally, the writing of video data as a product of the rendering algorithm is uniform, thereby improving the apparent video stability of the display to the user.

[0053] The client architecture 40 preferably executes in multiple process threads, with additional threads being utilized for individual network data request transactions. As shown in Figure 5, an image parcel management process 80 implements a loop that determines image parcels subject to update 82 and creates corresponding image parcel download requests 84. Navigation events that alter the viewing frustum are considered in part to determine the current field of view. The quad-tree data structures are examined 86 to identify viewable image parcels of higher resolution than currently available in the parcel data store 46.

[0054] A pool of image request threads is preferably utilized to manage the image parcel download operations. In the preferred embodiments of the present invention, a pool of four network request threads is utilized. The number of pool threads is determined as a balance between the available system resources and the network response latency, given the available bandwidth of the network connection. Empirically, for many wireless devices, four concurrent threads are able to support a relatively continuous delivery of image data parcels to the client 20 for display processing. As image parcels are progressively identified for

download, a free request thread is employed to issue 88 a corresponding network request to the server 12, 22. When a network response is received, the corresponding thread recovers 90 the image parcel data. The received image parcel is then stored 92 in a corresponding quad-tree data structure node.

[0055] For small clients 20, the available memory for the parcel data store 46 is generally quite restricted. In order to make optimal use of the available memory, only currently viewable image parcels are subject to download. Where the size of the parcel data store 46 is not so restricted, this constraint can be relaxed. In either case, a memory management process 94 runs to monitor use of the parcel data store 46 and selectively remove image parcels to free memory for newly requested image parcels. Preferably, the memory management process 94 operates to preferentially remove image parcels that are the furthest from the current viewing frustum and that have the highest data structure depth. Child node image parcels are always removed before a parent node parcel is removed.

[0056] A preferred network request management process 100 is shown in Figure 6. The process 100 waits 102 on the existence of a download request in the priority request queue 52. The process 100 then waits on a network request pool thread to become free 104. When a network request thread becomes available, the process 100 examines 106 all of the pending requests in the priority request queue 52 and selects 108 the request with the highest assigned priority. Thus, sequentially enqueued requests can be selectively issued out of order based on an independently assigned request priority. The request is then issued 110 and the request management process 100 leaves the request thread waiting on a network response.



[0057] Figure 7 presents a preferred display management process 120. Event driven user navigation information is evaluated 122 to determine a current viewing frustum location and orientation within a three-dimensional space relative to the displayed image. An algorithmic priority selection 124 of a next image parcel to render is then performed. The selected image parcel is then rendered 126 to the display memory 70. The rendering operation preferably performs a texture map transform of the parcel data corresponding to the current viewing frustum location and orientation. The overlay data is then parsed or is pre-parsed to determine 128 whether the image coordinates of any overlay annotation correspond to the current image parcel location. If the coordinates match, the overlay annotation is rendered 130 to the video display memory 70. The process 120 then continues with the next selection 124 of an image parcel to render, subject to any change in the viewing frustum location and orientation.

[0058] A preferred implementation of the selection 124 and rendering 126 of image parcels in accordance with the present invention is detailed in Figures 8 through 10. Referring first to Figure 8, any outstanding requests in the priority request queue 52 are preferably cleared 142 in response to a change in the viewing frustum location and orientation. The effective altitude of the viewing frustum and the resolution of the client display are then used as a basis for determining an optimal level of detail  $L$  that will be displayed. The detail level  $L$  value operates as a floor defining the maximum resolution  $K_L$  of image data that can be effectively viewed on the client display given the location and orientation of the viewing frustum. Constraining image parcel requests to the resolution range  $K_N$  to  $K_L$ , where  $K_N$  is the lowest resolution derivative image stored by the

network server 12, 22, prevents the download and processing of image parcels that cannot provide any perceptible improvement in the displayed image.

[0059] As part of the recursive evaluation of the optimal level of detail L, the image display space is progressively split 146 by four to one reductions into polygons. The quad-tree data structures holding existing image parcel data in the parcel data store 46 are concurrently traced 148 to establish a correspondence with the polygon map. Where the trace of a quad-tree data structure completes 150 to a node index of L for a polygon P, the node corresponding image parcel is associated with polygon P. The polygon P will not be further subdivided and no higher resolution image parcels will be requested for any portion of the image within the area represented by polygon P. Where the trace reaches a maximum node index of D for a polygon P' 152, where  $N \leq D < L$  and N is the index of the lowest resolution derivative image stored by the network server 12, 22, the image parcel associated with the node is associated with the polygon P'. This polygon P' will be subject to further subdivision and progressive requests for image parcels of higher resolution up to the detail level L.

[0060] Referring now to Figure 9, a display image is then rendered 160 beginning with the maximum depth polygons previously found. Iterating over the set of maximum depth polygons, any polygons outside of the viewing frustum are skipped 162. Polygons that are at least partially visible are clipped to the applicable bounds of the viewing frustum 164. The polygon corresponding image parcel data is then texture mapped 166 into the polygon corresponding coordinates of the video memory 70. If the node index depth of the rendered image parcel is at least equal to the prior determined optimal detail level L 168, the iteration over the polygons P continues.

[0061] Where the node index depth is less than the optimal detail level L 170, the polygon P' is subdivided into four polygons and correspondingly represented by the creation of four child nodes within the associated quad-tree data structure 172. Four image parcel download requests are then created 174.

[0062] The download priority associated with each request is determined 176 by execution of a function S that operates on a 2D polygon argument P and returns a real number representing the request priority. The function argument P is a list of real (x, y) coordinates of the vertices of the current polygon in screen coordinates after being clipped to fit within the current viewing frustum. That is, the function S works over general polygons in a two-dimensional space, whose vertices are specified by the series  $\{(x(1),y(1)),(x(2),y(2)),\dots,(x(n),y(n))\}$ . The argument P vertices sent to S represent the position of the vertices composing each of the polygons, after being clipping to the viewing frustum, viewable within the display space having the fixed resolution [xRes, yRes]. Thus, the clipped polygons are all within the rectangle  $[0, xRes] \times [0, yRes]$ .

[0063] In execution of the function S, each of the P coordinates is first transformed by linear mapping of the screen coordinate space to the square  $[-1,1] \times [-1,1]$  by the operation  $x(i) := (x(i) - xRes/2) / (xRes/2)$ ;  $y(i) := (y(i) - yRes/2) / (yRes/2)$ . The x and y coordinate values of each vertex  $(x(i),y(i))$  for  $i = 1$  to  $n$  are then transformed by the function  $T(a) = \text{sgn}(a) * \text{pow}(|a|, d)$ , where the control parameter d is a constant in the range  $(0,1]$ , or equivalently the interval  $0 < d \leq 1$ . The function S then returns a real value that is equal to the area covered by the argument polygon P vertices subject to the applied coordinate transformation. Thus, the accumulated priority for any image parcel pending download is the sum of the values of returned by the function S for each of the

viewable polygons that require some part of the image parcel as the source data for texture map rendering of the polygon. The priority operation of the request queue 52 is such that download requests will be issued preferentially for image parcels with the largest priority value.

[0064] In accordance with the preferred embodiments of the present invention, the value of the control parameter  $d$  can be adjusted to ultimately affect the behavior of the function  $S$  in determining the download request priority. In general, image parcels with lower resolution levels will accumulate greater priority values due to the larger number of polygons that may use a given low resolution image parcel as a rendering data source. Such lower resolution image parcels are therefore more likely to be preferentially downloaded. In accordance with the present invention, this generally assures that a complete image of at least low resolution will be available for rendering.

[0065] The control parameter  $d$ , as applied in execution of the function  $S$ , well as the area distortion produced by the projection transform also influences the value returned by the function  $S$  such that relatively higher-resolution image parcels near the image view point will occasionally achieve a higher priority than relatively remote and partially viewed image parcels of lower resolution. Using values smaller than 1 for the control parameter  $d$  results in requests with a higher priority for parcels covering areas near the focal point of the viewer, which is presumed to be the center point of the display space, relative to requests for parcels further from the center point in absolute terms and of the same resolution depth  $D$ . Thus, in accordance with the present invention, the priority assigned to image parcel requests effectively influences the order of requests based on the relative contribution of the image parcel data to the total display quality of the

image. Empirically, a value of 0.35 for the control parameter  $d$  for small screen devices, such as PDAs and webphones has been found to produce desirable results.

[0066] The computed priorities of each of the four newly created image parcel requests are then assigned 178 and the requests are enqueued in the priority request queue 52. The next polygon  $P$  is then considered in the loop of the image parcel rendering process 160.

[0067] The preferred algorithm 180 for determining the detail level  $L$  value for a given viewing frustum is shown in Figure 10. In accordance with the present invention, the optimal detail level  $L$  is effectively the limit at which the resolution of image parcel data functionally exceeds the resolution of the client display. Preferably, to determine the optimal detail level  $L$ , the viewpoint or camera position of the viewing frustum is determined 182 relative to the displayed image. A nearest polygon  $P$  of depth  $D$  is then determined 184 from the effective altitude and attitude of the viewpoint. The nearest point  $A$  of the polygon  $P$  is then determined 186. The point  $A$  may be within the interior or an edge of the polygon  $P$ , though most likely be located at a vertex of the polygon  $P$ .

[0068] The optimum level of detail  $L$  at point  $A$  is then computed 188 as the base-4 logarithm of the number of pixels on the screen that would be covered by a single pixel from an image parcel of the lowest resolution  $K_N$  image, which is the quad-tree root image and corresponds to an image area covering the entire image map. The point  $A$  optimal detail level  $L$  is preferably computed analytically from the local value of the Jacobian of the projective transform used to transform the three dimensional image coordinate space to screen coordinates, evaluated at the point  $A$ .

[0069] Where the depth D of the polygon P is greater than the depth of the computed optimal level of detail L, the detail level L is taken as the optimal detail level L 190. Thus, through the process 140, an image parcel or corresponding section of the closest resolution image parcel associated with a parent node in the quad-tree data structure relative to the depth level L will be used as the texture for rendering the polygon P. Conversely, if the depth D is less than that of the optimal detail level L, the polygon P is effectively split into quadrants and the optimal level of detail is reevaluated. The process 180 thus continues iteratively until the optimal detail level L is found.

[0070] Thus, a system and methods of optimally presenting image data on client systems with potentially limited processing performance, resources, and communications bandwidth have been described. While the present invention has been described particularly with reference to the communications and display of geographic image data, the present invention is equally applicable to the efficient communications and display of other high resolution information.

[0071] In view of the above description of the preferred embodiments of the present invention, many modifications and variations of the disclosed embodiments will be readily appreciated by those of skill in the art. It is therefore to be understood that, within the scope of the appended claims, the invention may be practiced otherwise than as specifically described above.